

Bayesian Model Selection of Stochastic Block Models

Xiaoran Yan

Indiana University Network Science Institute

Bloomington, Indiana 47408

Email: xiaoran.a.yan@gmail.com

Abstract—A central problem in analyzing networks is partitioning them into *modules* or *communities*. One of the best tools for this is the *stochastic block model*, which clusters vertices into blocks with statistically homogeneous pattern of links. Despite its flexibility and popularity, there has been a lack of principled statistical model selection criteria for the stochastic block model. Here we propose a Bayesian framework for choosing the number of blocks as well as comparing it to the more elaborate degree-corrected block models, ultimately leading to a universal model selection framework capable of comparing multiple modeling combinations. We will also investigate its connection to the minimum description length principle.

I. INTRODUCTION

An important task in network analysis is community detection, or finding groups of similar vertices which can then be analyzed separately [1]. Community structures offer clues to the processes which generated the graph, on scales ranging from face-to-face social interaction [2] through social-media communications [3] to the organization of food webs [4]. However, previous work often defines a “community” as a group of vertices with high density of connections within the group and a low density of connections to the rest of the network. While this type of *assortative* community structure is generally the case in social networks, we are interested in a more general definition of *functional* community—a group of vertices that connect to the rest of the network in similar ways. A set of similar predators form a functional group in a food web, not because they eat each other, but because they feed on similar prey. In English, nouns often follow adjectives, but seldom follow other nouns.

The *stochastic block model* (SBM) is a popular network model for such functional communities [5]. It splits vertices into latent blocks, within which all vertices are stochastically equivalent in terms of how they connect to the rest of the network [6]. As a generative model, it has a well-defined likelihood function with consistent parameter estimates. It also provides a flexible base for more general latent state models. In particular, Karrer and Newman proposed a variant called the degree-corrected block models (DC-SBM) capable of capturing specific degree distributions and Yaojia et al. generalized it to directed graphs [7], [8]. There are also other variants capable of modeling overlapping, hierarchical and even “meta-data” dependent communities [9], [10], [11].

Performance of different latent state models vary under different scenarios. For SBM and its generalizations, picking

the right model (*model selection*), and in particular picking the right number of blocks (*order selection*) is crucial for successful network modeling. Numerous statistical model selection techniques have been developed for classic independent data. Unfortunately, it has been a common mistake to use these techniques in network models without rigorous examinations of their compatibility, ignoring the fact that some of the fundamental assumptions has been violated by moving into the domain of relational data [12], [13], [14]. As a result, some employed these information criteria directly with out knowing the consequences [4], [15], while others remain skeptical and use it only when no alternatives is available [9], [16].

Our main contribution in this paper is to develop a Bayesian model selection framework for comparing multiple SBM variants with different number of blocks¹. In Section 2, we will first go over the Bayesian model selection approaches and the minimum description length (MDL) principle [17], [18]. In Section 3, we will propose Bayesian order selection for the SBM, leading to a Bayesian Information criterion (BIC) [19]. We will also establish the equivalence between BIC and MDL. In section 4, we will generalize these results to the DC-SBM, leading to a universal model selection framework capable of comparing multiple models that combines different model and order choices. We will compare its theoretic and empirical results to previous MDL based approaches [18], [20] as well as previous work on frequentist model selection [21].

II. BACKGROUND AND RELATED WORK

Model selection is about balancing the trade-off between model complexity and fit to the data. Models with more parameters have a natural advantage at fitting data. Simpler models have lower variability, and are less sensitive to noise in the data. A good model choice should avoid both over-fitting and under-fitting, and only include additional parameters when they do capture meaningful structures in the data [19], [22].

The frequentist approach to model selection cast the problem as a hypothesis testing. It focus on estimating the likelihood-ratio between a pair of candidate models under the null hypothesis. Such frequentist method has been used for

¹We will focus on the model/order selection problems of the SBM and Degree-corrected block model. For mathematical convenience, we will define these models for undirected graphs. Directed versions require additional specifications but the general form remains the same [8]. Readers should be able to generalize the result in this paper to these cases.

both model selection and order selection [21], [23]. In this paper, we shall follow the other major school of statistics and use Bayesian techniques for model selection.

A. Bayesian model selection

While point estimates like the maximum likelihood lead to over-fitting, Bayesian approaches take the whole posterior distribution into account, thus achieving the trade-off between model complexity and its fit to data [19], [22]. These posteriors distributions can be formulated using Bayes' rule,

$$P(M_i | G) = \frac{P(M_i)}{P(G)} P(G | M_i) \propto \iiint_0^1 P(G | M_i, \Pi_i) P(\Pi_i | M_i) d\Pi_i, \quad (1)$$

where we have assumed a uniform prior of models $P(M_i)$, and the total evidence of data $P(G)$ is constant for all models.

The posterior $P(M_i | G)$ has an intuitive interpretation for model selection. It is proportional to $P(G | M_i)$, which is the integrated complete likelihood (as ICL in [24]) $P(G | M_i, \Pi_i)$ over the prior of parameters $P(\Pi_i | M_i)$. To compare models, the standard approach is to divide one posterior with another in a likelihood-ratio style, leading to the Bayes factor [25]. Unlike the frequentist tests, Bayes factor uses the ICL without any dependence on parameters. Therefore, it can be applied not only to nested model pairs, but two models of any form.

Ironically, instead of one preferred choice, a fully Bayesian approach would give a posterior distribution over all candidate models. The maximum a posteriori (MAP) method presented in this paper is technically a hybrid between Bayesian and frequentist ideas. For comparison with frequentist and MDL based methods, we will still call it Bayesian. Nonetheless, most results can be adapted for full Bayesian analysis.

B. Bayesian information criterion

While Bayesian model selection is compatible for models of any form, the posteriors can often be intractable. The exception is when the likelihood model is from a family with conjugate priors. The posteriors will then have closed form solutions. Fortunately for us, the block models fall under this category.

Bayes factor is also quite cumbersome when more than two candidate models are involved. Instead of the principled hypothesis testing framework, Bayesian information criterion (BIC) [26] gives standard values for each candidate model:

$$BIC(M_i) = -2 \ln P(Y | M_i, \hat{\Pi}_i) + |\Pi_i| \ln n, \quad (2)$$

where $|\Pi_i|$ is the degree of freedom of the model M_i with a parameter set $|\Pi_i|$, and n is number of i.i.d. samples in the data. The above simple formulation consists of a maximized likelihood term and a penalty term for model complexity, intuitively corresponding to the trade-off we are looking for. As we will later show, it is in fact a large sample approximation to twice the logarithm of the ICL (Equation (1)).

BIC has been applied to different clustering models of i.i.d. data [27], [24], [28]. Handcock et al. proposed a variants of BIC for order selection of a latent space clustering model on

networks [15]. Recently, Côme and Latouche derived a BIC from the ICL for the vanilla SBMs [29]. We will redo the derivation in our notions and generalize it to degree-corrected block models in this paper.

C. The minimum description length principle

By compressing data with different coding schemes, information theory has a long history dealing with the trade-off between complexity and fit. Searching for the model with best predictive performance is essentially finding a coding scheme that lead to the minimum description length (MDL) [17], [30]. Under the MDL principle, the trade-off takes the form of balancing between the description length of the coding scheme and that of the message body given the code [18], [31].

MDL is closely related with Bayesian model selection, particularly the BIC formulation [17]. In [20], [32], Peixoto demonstrated that for basic SBMs, the ICL or Bayesian posteriors are mathematically equivalent to MDL criteria under certain model constrains. This equivalence, as we will show in this paper, underlies a fundamental connection in the form of carefully designed Bayesian codes, which can be derived from the ICL (1) with intuitions.

III. BAYESIAN ORDER SELECTION OF THE SBM

In this section, we will formally introduce the SBM, derive the ICL for the SBM order selection problem, and design an intuitive Bayesian code to demonstrate the fundamental connection between Bayesian order selection and the MDL. Finally we will propose a Bayesian Information criterion (BIC) [19] for order selection of the SBM.

A. The stochastic block model

We represent our network as a simple undirected graph $G = (V, E)$, without self-loops. G has n vertices in the set V , m edges in the set E , and they can be specified by an adjacency matrix A where each entry $A_{uv} = 0$ or 1 indicates if there is an edge in-between. We assume that there are k blocks of vertices (choosing k is the order selection problem), so that each vertex u has a block label $g(u) \in \{1, \dots, k\}$. Here $n_s = |\{u \in V : g(u) = s\}|$ is the number of vertices in block s , and $m_{st} = |\{u < v, (u, v) \in E : g(u) = s, g(v) = t\}|$ is the number of edges connecting between block s and block t .

We assume that G is generated by a SBM, or a ‘‘vanilla SBM’’ as we will call it throughout this paper for distinction. For each pair of vertices u, v , there is an edge between u and v with the probability $p_{g(u), g(v)}$ specified by the $k \times k$ block affinity matrix p . Each vertex label $g(u)$ is first independently generated according to the prior probability $q_{g(u)}$ with $\sum_{s=0}^k q_s = 1$. Given a block assignment, i.e., a function $g : V \rightarrow \{1, \dots, k\}$ assigning a label to each vertex, the probability of generating a given graph G in this model is

$$P(G, g | q, p) = \prod_{s=1}^k q_s^{n_s} \prod_{s < t} p_{st}^{m_{st}} (1 - p_{st})^{n_s n_t - m_{st}}, \quad (3)$$

This likelihood factors into terms for vertices and edges, conditioned on their parameters q, p respectively.

Take the log of Equation (3), we have the log-likelihood

$$\begin{aligned} \log P(G, g | q, p) &= \sum_{s=1}^k n_s \log q_s \\ &+ \sum_{s \leq t}^k (m_{st} \log p_{st} + (n_s n_t - m_{st}) \log(1 - p_{st})) . \end{aligned} \quad (4)$$

B. Bayesian posterior of the SBM

A key design choice of MAP model selection is picking which parameters to integrate. By being partially frequentist, we have the freedom to decide how "Bayesian" we wanted to be. The decision can also be understood as balancing between the bias in the learning task and the variance in application domains. For example, if the learning task is to find the latent state model with the most likely parametrization, regardless of specific latent block assignments, the integral over parameters is not necessary, as in the paper [21].

Alternatively, if the learning task is to find the model with the most likely latent state, we do not need the sum over the latent state g , and benefit from the smaller variance because of the bias we are willing to assume. However, if we plan to apply the learned latent state model to similar networks with different parameterizations, the integral over parameters remains essential. This corresponds to the idea of *Universal Coding* in information coding theories [17], where a code has to achieve optimal average compression for all data generated by the same code scheme without knowing its parametrization a priori. This is the approach Handcock et al. adopted for order selection of their latent space model [15]. The same method was used in the active learning algorithm for vanilla SBMs [33]. In this paper, we will follow this approach for its connection with MDL methods. We will use a Monte Carlo sampling method to find the most likely latent state \hat{g} .

According to equation (1), we have the posterior of a SBM M_i with the parameters $\{p, q\}$:

$$\begin{aligned} P(M_i | G) &= \frac{P(M_i)}{P(G)} P(G | M_i) \\ &\propto \sum_g \iiint_0^1 d\{p_{st}\} d\{q_s\} P(G, g | p, q) , \end{aligned} \quad (5)$$

where we assume the prior of models $P(M_i)$ is uniform, and the total evidence of data $P(G)$ is constant. Here the ICL $P(G | M_i)$ is integrated over both p and q entries, as well as summed over all latent states g . Since we are interested in the most likely latent state \hat{g} , we can forgo the sum,

$$P(G, g | M_i) = \iiint_0^1 d\{p_{st}\} d\{q_s\} P(G, g | p, q) . \quad (6)$$

with $P(G, \hat{g} | M_i) = \max_g P(G, g | M_i)$.

If we assume that the p_{st} and q_s entries are independent, conditioned on the constrain $\sum_s q_s = 1$, and they follows their

respective conjugate Dirichlet and Beta priors (δ and $\{\alpha, \beta\}$ respectively), we have

$$\begin{aligned} P(M_i | G) &\propto \iiint_0^1 d\{p_{st}\} d\{q_s\} P(G, g | q, p) \\ &= \left(\int_{\Delta} dq \text{Dirichlet}(\vec{q} | \vec{\delta}) \prod_{s=1}^k q_s^{n_s} \right) \\ &\quad \left(\prod_{s \leq t}^k \int_0^1 dp_{st} \text{Beta}(p_{st} | \alpha, \beta) p_{st}^{m_{st}} (1 - p_{st})^{n_s n_t - m_{st}} \right) \\ &= \left(\frac{\Gamma(\sum_{s=1}^k \delta_s) \prod_{s=1}^k \Gamma(n_s + \delta_s)}{\prod_{s=1}^k \Gamma(\delta_s) \Gamma(\sum_{s=1}^k (n_s + \delta_s))} \right) \\ &\quad \left(\prod_{s \leq t}^k \frac{\Gamma(\alpha + \beta) \Gamma(m_{st} + \alpha) \Gamma(n_s n_t - m_{st} + \beta)}{\Gamma(\alpha) \Gamma(\beta) \Gamma(n_s n_t + \alpha + \beta)} \right) \\ &= P(V, g | M_i) \times P(E, g | M_i) , \end{aligned} \quad (7)$$

where we have assumed the same beta prior $\{\alpha, \beta\}$ for all p_{st} entries, and applied the Euler integral of the first kind, and its multinomial generalization on the simplex $\sum_s q_s = 1$.

Equation (7) shows, the ICL factors into terms for vertices and edges. It holds for any latent state g , as long as n_s and m_{st} terms are consistent with the given g . For empirical verifications, we will use a Monte Carlo sampling method to find the most likely latent state \hat{g} .

To get an idea of the posterior distribution $P(M_i | G)$, we assume that the data G follows a uniform prior over random graphs generated by a SBM with 5 prescribed blocks. For simplicity, we have also plugged in the uniform priors (i.e., $\delta_{v_s} = 1, \alpha = \beta = 1$) for the parameters, as it is done in [33],

$$\begin{aligned} P(M_i | G, g) &\propto P(G, g | M_i) \\ &= \left((k-1)! \frac{\prod_{s=1}^k n_s!}{(n+k-1)!} \right) \left(\prod_{s \leq t}^k \frac{m_{st}! (n_s n_t - m_{st})!}{(n_s n_t + 1)!} \right) . \end{aligned} \quad (8)$$

The distributions of the posterior with different number of blocks are shown in Figure 1 (top). While the SBM with correct number of blocks (red) does has slightly higher likelihood in average, it overlaps quite heavily with SBMs with fewer blocks (green) or more blocks (blue). Further investigation reveals that most of the variance came from the randomness in the generated data. Once we fix the input graph for all the candidate models, the SBM with correct number of blocks always has a higher likelihood than the others, as illustrated in Figure. 1 (bottom).

C. Bayesian code for the SBM

Now we will design a Bayesian code based on the ICL. According to Grünwald [17], the equivalence between Bayesian and the MDL principle for model selection holds in general. In our case, if we choose the Jeffreys priors for p_{st} and q_s entries (i.e., $\alpha = \beta = \delta_{v_s} = 1/2$), the coding length according to the Bayesian model is asymptotically the same as the optimal

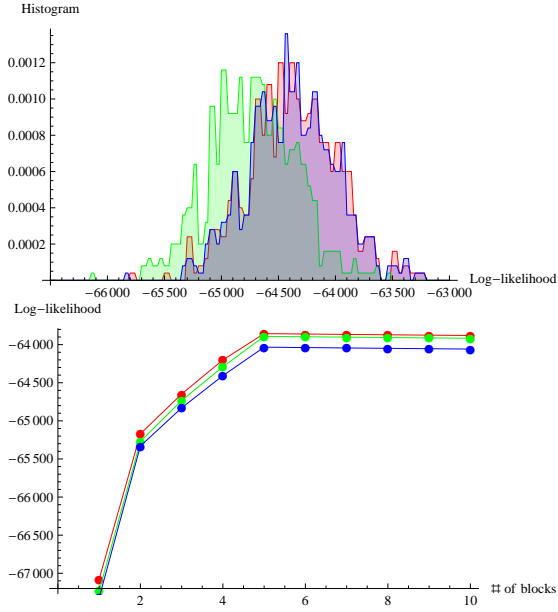


Fig. 1. Top: the histogram of log-integrated complete likelihoods (log-ICL) of the Bayesian model (Equation (8)). The distributions are gathered from randomly generated SBMs with 1000 vertices and 5 prescribed blocks. SBMs with different number of blocks k are fitted to the data. Specifically, the green distribution is from a SBM with $k = 4$, the red with $k = 5$ and the blue with $k = 9$. The experiment is done using Monte Carlo sampling, with 500 total samples of G . Bottom: the change of log-ICL for sample graphs as the number of blocks grows. Three graphs are randomly generated according to the same prescribed blocks. All of them have highest log-likelihood at the correct number of blocks $k = 5$.

minimax coding. Grünwald also pointed out in [17], while the Jeffreys priors lead to the optimal universal coding, other priors and their corresponding non-optimal coding still lead to description length of the same asymptotic order, if the prior is dominated by the evidence.

This justifies the uniform prior assumptions we used in Equation (8), and we can rewrite it as

$$P(G, g | M_i) = P(V, g | M_i) \times P(E, g | M_i) = \left(\frac{1}{\binom{n+k-1}{k-1}} \frac{1}{\binom{n}{(n_1, n_2, \dots, n_k)}} \right) \left(\prod_{s \leq t}^k \frac{1}{\binom{n_s n_t}{m_{st}} (n_s n_t + 1)} \right). \quad (9)$$

The dominating terms in Equation (9) lead to a Bayesian universal code for a graph G consists of the following parts:

- 1) number of blocks k ($\log k$ bits, implicit)
- 2) code for the partition of n into k n_s terms ($\log \binom{n+k-1}{k-1}$ bits,), which will specify a block size sequence (ordered in terms of blocks)
- 3) code for assigning each vertex to blocks according to the n_s terms ($\log \binom{n}{(n_1, n_2, \dots, n_k)}$ bits)
- 4) for each pair of blocks s, t , the number of undirected edges m_{st} going between block s and block t ($\log m_{st} < \log(n_s n_t + 1)$ bits)

- 5) for each pair of blocks s, t , code for the edge allocations given m_{st} ($\log \binom{n_s n_t}{m_{st}}$ bits, uniformly random allocations)

According to [17], a mapping exists between probability distributions and prefix codes. In the Bayesian code, the distribution of possible realizations in part i ($i > 1$) conditioned on all previous code parts are all uniform, the optimal code length for part i thus can be quantified by the negative logarithm of the corresponding combinatorial terms in Equation (9).

The aforementioned Bayesian code is identical to the description length for the single level model discussed in [20]. In fact, Peixoto mathematically arrived at the same equivalence in the appendix of [20]. Earlier formulations of MDLs, however, are usually defined in terms of the entropy minimizers of the likelihood functions, which is equivalent to the maximum likelihood formulation in the BIC (Equation (2)).

D. BIC for order selection of SBMs

The key to transform Equation (7) to the BIC formulation (Equation (2)) is the Laplace's approximation with uniform priors. Or equivalently, by using Stirling's formula on the factorials in Equation (9),

$$\begin{aligned} P(G, g | M_i) &= P(V, g | M_i) \times P(E, g | M_i) \\ &\approx \frac{\prod_{s=1}^k \sqrt{2\pi n_s}}{\binom{n+k-1}{n}} \prod_u \frac{n_{g(u)}}{\sqrt{2\pi n}} \prod_{s \leq t}^k \frac{2\pi \sqrt{m_{st}(n_s n_t - m_{st})}}{\sqrt{2\pi (n_s n_t)} (n_s n_t + 1)} \\ &\quad \prod_{u < v, (u,v) \in E} \frac{m_{g(u)g(v)}}{n_{g(u)} n_{g(v)}} \prod_{u < v, (u,v) \notin E} \left(1 - \frac{m_{g(u)g(v)}}{n_{g(u)} n_{g(v)}} \right) \\ &\approx P(V, g | \hat{q}) \frac{\prod_{s=1}^k \sqrt{2\pi n_s}}{\binom{n+k-1}{n} \sqrt{2\pi n}} \\ &\quad \times P(E, g | \hat{p}) \prod_{s \leq t}^k \frac{\sqrt{2\pi}}{\sqrt{\frac{n_s^3 n_t^3}{m_{st}(n_s n_t - m_{st})}}}, \end{aligned} \quad (10)$$

where we plugged in the MLEs $\hat{q}_s = \frac{n_s}{n}$ and $\hat{p}_{st} = \frac{m_{st}}{n_s n_t}$.

If we take the negative log of Equation (10). The factor associated with E becomes:

$$\begin{aligned} &-\ln P(E, g | M_i) \\ &\approx -\ln P(E, g | \hat{p}) - \sum_{s \leq t}^k \frac{1}{2} \ln \frac{2\pi m_{st}(n_s n_t - m_{st})}{n_s^3 n_t^3} \\ &\approx -\ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \frac{n^6}{4\pi |E| (n^2 - |E|)} - C, \end{aligned} \quad (11)$$

where we made a mean-field assumption about both n_s and m_{st} under constant number of blocks k . If the edge density of graph scales as $|E| = \rho n^2$, with ρ being a constant such that $0 \leq \rho \ll 1$, we have

$$\begin{aligned} -\ln P(E, g | M_i) &\approx -\ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \frac{n^6}{4\pi \rho n^4 (1 - \rho)} \\ &\approx -\ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \Theta(n^2). \end{aligned}$$

Putting it together with the term associated with V , in which we again assumed mean-field n_s terms, we get

$$\begin{aligned} -\ln P(G, g | M_i) &= -\ln P(V, g | M_i) - \ln P(E, g | M_i) \\ &\approx -\ln P(V, g | \hat{q}) + \Theta(k \ln n) - \ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \Theta(n^2) \\ &= -\ln P(G, g | \hat{p}, \hat{q}) + \frac{k^2}{2} \ln \Theta(n^2). \end{aligned} \quad (12)$$

Multiply by 2, we have the BIC for order selection in SBMs:

$$BIC_{SBM}(M_i) = -2 \ln P(G, g | M_i, \hat{\Pi}_i) + k^2 \ln \Theta(n^2), \quad (13)$$

with k^2 specifying the number of parameters in the block affinity matrix p and n^2 represent the sample size as pairwise edge/non-edge interactions.

(13) is simply the direct application of BIC to the SBM as it is defined in Equation (3). In [15], Handcock et al. arrived at the same equation without showing derivations. They also suggested using $|E|$ instead of n^2 as the sample size measure, on the basis of arguments that apply to a very different model. However, the above derivation is no longer correct when the edge density scales differently. If the graph is sparser with $|E| = \rho n$, we have

$$\begin{aligned} -\ln P(E, g | M_i) &\approx -\ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \frac{n^6}{4\pi\rho n^2(n-\rho)} \\ &\approx -\ln P(E, g | \hat{p}) + \frac{k^2}{2} \ln \Theta(n^3). \end{aligned} \quad (14)$$

The sparse BIC for order selection is then:

$$BIC_{SBM}^*(M_i) = -2 \ln P(G, g | M_i, \hat{\Pi}_i) + k^2 \ln \Theta(n^3), \quad (15)$$

where the penalty term becomes even greater, favoring simpler models to compensate for sparser data samples (edges).

IV. BAYESIAN MODEL SELECTION OF THE DC-SBM

In this section, we will generalize the results to the degree-corrected block model, ultimately leading to a universal model selection framework capable of comparing multiple models with different number of blocks.

A. The Degree-corrected block model

The vanilla SBM assumes that each entry A_{uv} is 0 or 1. Another restriction of the vanilla SBM is that all the vertices in the same block have the same expected degree, following a Binomial distribution with a narrow peak. As a consequence, it ‘‘resists’’ putting vertices with very different degrees in the same block, leading to problems in real networks when the degree distribution is heavy tailed.

The DC-SBM addresses these problem by allowing degree heterogeneity within blocks. Each vertex gets an additional parameter θ_u , which scales the expected number of edges connecting to it [7]. DC-SBM also generalizes the edge generating process to Poisson, thus allowing multi-edges between vertices. According to the block assignment g , the means of these Poisson draws depends on the $k \times k$ block affinity matrix

ω , which replaces the p matrix in the vanilla SBM. The edge generating likelihood is now

$$A_{uv} | g \sim \text{Poi}(\theta_u \theta_v \omega_{g_u g_v}).$$

The parameter θ_u gives us control over the expected degree of each vertex, which for instance, could be a measure of popularity in social networks. The likelihood stays the same if we scale θ_u for all vertices in block s , provided we also divide ω_{st} for all s by the same factor. Thus identification demands an additional constraint. Here we use a convenient one that forces θ_u to sum to the total number of vertices in each block: $\sum_{u: g_u = s} \theta_u = n_s$. The ICL of DC-SBM is then

$$\begin{aligned} P(G, g | \theta, \omega, q) &= \prod_u q_{g_u} \prod_{u < v} \frac{(\theta_u \theta_v \omega_{g_u g_v})^{A_{uv}}}{A_{uv}!} \exp(-\theta_u \theta_v \omega_{g_u g_v}) \\ &= \prod_u \theta_u^{d_u} \prod_{s=1}^k q_s^{n_s} \prod_{s \leq t}^k \omega_{st}^{m_{st}} \exp(-n_s n_t \omega_{st}) \prod_{u < v} \frac{1}{A_{uv}!} \quad (16) \\ &= P(\Theta, g | \theta) \times P(V, g | q) \times P(E, g | p), \end{aligned}$$

where d_u is the degree of vertex u , and $P(\Theta, g | \theta)$ is the only factor containing the θ parameters.

DC-SBM can be simplified when modeling simple graphs. The last term becomes 1, and if we take the logarithm,

$$\begin{aligned} \log P(G, g | \theta, \omega, q) &= \sum_u d_u \log \theta_u + \sum_{s=1}^k n_s \log q_s + \sum_{s \leq t}^k (m_{st} \log \omega_{st} - n_s n_t \omega_{st}). \end{aligned}$$

Compare it with Equation (4), if the graph is not very dense such that $m_{st} \ll n_s n_t$,

$$\begin{aligned} \log P(G, g | q, p) &= \log P(G, g | 1, \omega, q) \\ &= \sum_{s=1}^k n_s \log q_s + \sum_{s \leq t}^k (m_{st} \log p_{st} - (n_s n_t - m_{st}) \log(1 - p_{st})). \end{aligned}$$

In other words, when $\omega_{st} \approx p_{st}$ and both approach 0, multi-edges are so rare that the DC-SBM becomes the vanilla SBM by setting $\theta_u = 1$ for all u . This nested model relationship is consistent with their Poisson counterparts in [21]. For mathematical convenience, we shall automatically make these approximations in the following sections².

B. Bayesian posterior of the DC-SBM

We will now generalize the Bayesian framework to the DC-SBM (Equation (16)). By Bayes’ theorem, we have the posterior of a DC-SBM M_i with the parameters $\{\theta, \omega, q\}$:

$$\begin{aligned} P_{DC}(M_i | G) &= \frac{P(M_i)}{P(G)} P_{DC}(G | M_i) \\ &\propto \int \int \int_0^1 d\{\theta_u\} d\{\omega_{st}\} d\{q_s\} P(G, g | \theta, \omega, q), \end{aligned} \quad (17)$$

²Notice that it is different from the notion of sparsity in Equation (15). Even as $\omega_{st} \approx p_{st} \rightarrow 0$, we can still have quadratic scaling of edge densities.

where we again assumed a uniform prior over $P(M_i)$, and a constant $P(G)$.

If θ_u , w_{st} and q_s entries are independent, with the constrains $\sum_{u: g_u=s} \theta_u = n_s$ and $\sum_i q_s = 1$, we have the posterior:

$$\begin{aligned} P_{DC}(M_i | G, g) &\propto P_{DC}(G, g | M_i) \\ &= \int \int \int_0^1 d\{\theta_u\} d\{w_{st}\} d\{q_s\} P(G, g | \theta, \omega, q) \\ &= \left(\int_{\Delta} d\{\theta_u\} \prod_u \theta_u^{d_u} \right) \int \int \int_0^1 d\{w_{st}\} d\{q_s\} P(G, g | \omega, q) \\ &\approx P(\Theta, g | M_i) \times P(G, g | M_i). \end{aligned} \quad (18)$$

The integrated DC-SBM has one additional factor, forming a pair of nested models with the integrated vanilla SBM.

To prepare $P(\Theta, g | M_i)$ for Bayesian treatments, we first change the variables $\theta_u = n_{g(u)} \eta_u$ in the first integral, making the integrand a proper multinomial distribution. Now if the new parameters η_u follow their Dirichlet conjugate priors,

$$\begin{aligned} P(\Theta, g | M_i) &\approx \prod_{s=1}^k \left(\int_{\Delta} d\{\eta_u\} \prod_{g(u)=s} \eta_u^{d_u} \right) \times \prod_u n_{g(u)}^{d_u+1} \\ &= \prod_{s=1}^k \left(\int_{\Delta} d\eta \text{Dirichlet}(\vec{\eta}_s | \vec{\gamma}_s) \prod_{g(u)=s} \eta_u^{d_u} \right) \times \prod_u n_{g(u)}^{d_u+1} \\ &= \prod_{s=1}^k \left((n_s - 1)! \frac{\prod_{g(u)=s} d_u!}{(D_s + n_s - 1)!} \right) \times \prod_u n_{g(u)}^{d_u+1}, \end{aligned} \quad (19)$$

where we applied the multinomial Euler integral on the simplex $\sum_{u: g_u=s} \eta_u = 1$. At the last line of the derivation, we again assume the priors are uniform (i.e., $\gamma_{\forall u} = 1$).

Following the derivation of BIC for the vanilla SBM, we apply the Stirlings formula to the factorials in (19),

$$\begin{aligned} P(\Theta, g | M_i) &\approx P(\Theta, g | \hat{\eta}) \prod_{s=1}^k \left(\frac{\prod_{g(u)=s} \sqrt{2\pi d_u}}{((D_s + n_s - 1) \sqrt{2\pi D_s})} \right) \times \prod_{s=1}^k n_s^{D_s + n_s} \\ &= P(\Theta, g | \hat{\theta}) \prod_{s=1}^k \left(\frac{\prod_{g(u)=s} \sqrt{2\pi d_u}}{((D_s + n_s - 1) \sqrt{2\pi D_s})} \right) \times \prod_{s=1}^k n_s^{n_s}, \end{aligned} \quad (20)$$

where $D_s = \sum_{g(u)=s} d_u$ is the total degree of vertices in block s . We have also plugged in the MLEs $\hat{\eta}_u = \frac{\hat{\theta}_u}{n_{g(u)}} = \frac{d_u}{D_{g(u)}}$.

Putting back the factors from the vanilla SBM (Equation (12)), and take the logarithm of it, we have the log-ICL:

$$\begin{aligned} \ln P_{DC}(G, g | M_i) &= \ln P(\Theta, g | M_i) + \ln P(G, g | M_i) \\ &\approx \ln P(G, g | \hat{\theta}, \hat{q}, \hat{p}) - \frac{k^2}{2} \ln \Theta(n^*) \\ &\quad + \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) - \Theta\left(n \ln \frac{2|E|}{n}\right), \end{aligned} \quad (21)$$

with again mean-field assumptions about n_s and D_s terms. Notice that $\Theta(n^*)$ is a general form for the correct sample size for graphs with different edge density scalings, corresponding to both Equation (12) and (14).

The blue curve in Figure 2 shows the empirical results. Here data G is generated by a DC-SBM with $n = 1000$ vertices




Fig. 2. The change of log-likelihood for given graphs as the number of blocks grows. The graph is a randomly generated DC-SBM with 1000 vertices and 5 prescribed blocks. Within each block, the degrees follow a bimodal distribution comprised of two Poisson distributions with their means 3 times apart. Both the vanilla SBM (red) and DC-SBM (blue) with different number of blocks k are fitted to the data. The experiment is done using a Monte Carlo sampling method. The log-likelihood values shown here has been normalized for both models so that they are comparable across models. (The normalization method will be formally introduced in the next section.

from 5 prescribed blocks and the same expected number of total edges. Degrees within each block now follows a bimodal distribution. This degree heterogeneity forces the vanilla SBM to split vertices into separate high degree and low degree blocks, while the DC-SBM can comfortably mix them together in the same block. As a consequence, posterior of the DC-SBM achieves much higher log-ICL with fewer blocks. It also correctly captures the correct number of blocks at $k = 5$, unlike the monotonic increasing log-ICL of the vanilla SBM.

C. Towards a universal model selection framework

In theory, a Bayesian approach based on the full ICL of Equation (5) can be used for comparing multiple SBM variants with different number of blocks together. However, our partially Frequentist approaches come with additional complications. In the previous example (Figure. 2), we took an extra step to normalize the log-ICL for the vanilla SBM and DC-SBM so that they are comparable. The normalization step is required because our MAP approach leaves some parameters out of integration or summation. As a result, the ICL is still dependent of these parameters, eventually leading to divergence in maximum likelihoods for different models. In our case, these are the block assignment variables g , which can change drastically from the vanilla SBM to the DC-SBM given the same graph.

To remedy the situation, we propose a normalization method that use a hybrid frequentist-Bayesian method by calculating the expected difference between Equation (12)/(14) and (21):

$$\begin{aligned} &\ln P_{DC}(G, \hat{g} | M_i) - \ln P_{SBM}(G, \hat{g}' | M_i) \\ &\approx \ln P(G, \hat{g} | \hat{\theta}, \hat{q}, \hat{\omega}) - \ln P(G, \hat{g}' | 1, \hat{q}', \hat{p}) \\ &\quad + \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) - \Theta\left(n \ln \frac{2|E|}{n}\right), \end{aligned}$$

where \hat{g} and \hat{g}' are the most likely block assignments for the DC-SBM and vanilla SBM respectively.

Equation (18) shows that the vanilla SBM and the DC-SBM still forms a pair of nested models after the partial integrations. Therefore the analysis in [21] holds. If we assume the underlying data is generated by the simpler vanilla SBM, we have both models converge to the same values for shared parameters, i.e. $\hat{g} = \hat{g}'$, $\hat{q} = \hat{q}'$ and $\hat{\omega} = \hat{p}$. In [21], the authors came into the conclusion that the difference between the maximum likelihood under the null model roughly follows a χ^2 distribution with a degree of freedom $n - k$, but corrections are needed when the graph is sparse.

To verify that we have the same maximum log-likelihood ratio, we can rewrite (12) and (21) as:

$$\begin{aligned}\ln P(G, g | 1, \hat{q}, \hat{p}) &\approx \ln P_{SBM}(G, g | M_i) + \frac{k^2}{2} \ln \Theta(n^*), \\ \ln P(G, g | \hat{\theta}, \hat{q}, \hat{p}) &\approx \ln P_{DC}(G, g | M_i) + \frac{k^2}{2} \ln \Theta(n^*) \\ &\quad - \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) + \Theta\left(n \ln \frac{2|E|}{n}\right).\end{aligned}$$

Therefore, we have the log-likelihood ratio,

$$\begin{aligned}\Lambda_{DC}(G, g) &= \ln P(G, g | \hat{\theta}, \hat{q}, \hat{p}) - \ln P(G, g | 1, \hat{q}, \hat{p}) \\ &\approx \ln P(\Theta, g | M_i) - \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) + \Theta\left(n \ln \frac{2|E|}{n}\right) \\ &= \ln P(\Theta, g | \hat{\theta}),\end{aligned}$$

which is exactly the same as the log-likelihood ratio for hypothesis testing in paper [21]. The agreement between Bayesian and Frequentist methods is not a coincident, because we have used uniform priors in our derivation.

Now we are ready to use the result in [21] for estimating the expected difference between the log-ICLs,

$$\begin{aligned}\mathbb{E}[\ln P_{DC}(G, \hat{g} | M_i) - \ln P_{SBM}(G, \hat{g}' | M_i')] \\ = \mathbb{E}[\Lambda_{DC}(G, g)] + \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) - \Theta\left(n \ln \frac{2|E|}{n}\right) \\ \approx \ln\left[\left(\frac{1}{2} + \frac{n}{24|E|}\right)(n - k)\right] + \Theta\left(\frac{n}{2} \ln \frac{n}{k}\right) - \Theta\left(n \ln \frac{2|E|}{n}\right),\end{aligned}\quad (22)$$

where \bar{d}_s is the average degree of vertices in block s and $\frac{n}{24|E|}$ is the first order correction for sparse graphs [21].

Equation (22) only holds when the vanilla SBM is the generative model of the data. For datasets in general, however, we can guarantee this happen by making sure the number of blocks k is so large that even the vanilla SBM over-fits. Going back to Figure 2, the red curves peaks at $k = 10$. We can then normalized the log-likelihood of the DC-SBM by subtracting it with the expected difference (Equation (22)) at $k = 10$.

Subtracting Equation (22) from (21), we have the normalized posterior of the DC-SBM which is now directly comparable to Equation (12)/(14):

$$\begin{aligned}\ln P_{DC}(G, g | M_i) &= \ln P(G, g | \hat{\theta}, \hat{q}, \hat{p}) \\ &\quad - \frac{k^2}{2} \ln \Theta(n^*) - \ln\left[\left(\frac{1}{2} + \frac{n}{24|E|}\right)(n - k)\right].\end{aligned}\quad (23)$$

Normalization allows us to compare multiple models with different k together. Now we can ask questions like which model should we use conditioned on a given number of blocks. If we compare the two models for the same k according to Figure 2, the choice would be DC-SBM when $k < 3$, and the vanilla SBM when $k > 6$ (There is no clear winner when $3 \leq k \leq 6$). The best model overall is a vanilla SBM with 10 blocks, which is consistent with our generative model (5 blocks with bimodal degree distributions).

To arrive at the BIC formulation, we discard the constants in Equation (23) and multiply both sides by -2 ,

$$\begin{aligned}BIC_{DC}(M_i) &= -2 \ln P(G, g | M_i, \hat{\Pi}_i) \\ &\quad + k^2 \ln \Theta(n^*) + 2 \ln[\Theta(n)].\end{aligned}\quad (24)$$

Compared with the vanilla SBM (Equation (13)/(15)), BIC_{DC} has an additional penalty term which grows with the size of the network. Therefore, DC-SBM favors fewer blocks than the vanilla SBM, since the flexibility provided by the additional parameters allow vertices with very different degrees to coexist in the same block.

D. Results on real world networks

To illustrate how the universal framework might work on real world data sets, we investigate two simple social networks using Equation (13) and (23). The first is a social network consisting of 34 members of a karate club, where undirected edges represent friendships [2]. The network is made up of two assortative blocks, each with one high-degree hub and many low-degree peripheral vertices. In [21], the authors studied the model selection problem between the vanilla SBM and DC-SBM conditioned on $k = 2$. Using the frequentist likelihood ratio test, there were not enough evidence to reject the null hypothesis that the network is generated by a vanilla SBM.

The result based on the log-ICL with normalization confirms that the DC-SBM has a higher likelihood at $k = 2$ (see Figure 3, top). However, for DC-SBM, it maximizes at $k = 1$, which means with degree correction, any blocking leads to over-fitting. Therefore, for any meaningful communities, the better choice is the vanilla SBM with a bigger k , because the degree heterogeneity is not strong enough to justify the DC-SBM. In fact, the best model according to the universal Bayesian framework is the vanilla SBM with $k = 4$, which corresponds to the division with high/low degree blocks for each cluster. This is also reminiscent to the result using active learning [33], where the vanilla SBM labels most of the vertices correctly once the high degree vertices are known.

The second example is a network of political blogs in the US [3]. Here we focus on the largest component with 1222 blogs and 19087 links between them. The blogs have known political leanings, with either liberal or conservative labels. The network is assortative, with heavy tailed degree distributions within each block. As a consequence, the frequentist analysis in [21] supported the hypothesis that the network is generated by a DC-SBM. The universal Bayesian framework confirms the previous conjectures based on frequentist arguments at $k = 2$ (see Figure 3, bottom). In fact, the degree heterogeneity




Fig. 3. The change of log-ICL for the karate club network (top) and the political blogs network (bottom) as the number of blocks grows. Both the vanilla SBM (red) and DC-SBM (blue) with different k are fitted to the data. The experiment is done using a Monte Carlo sampling method. The log-likelihood values shown here has been normalized for both models.

here is so strong, that the vanilla SBM never overtake DC-SBM even with very big k values. If you are interested in large scale community structure of the political blogs network, such as the political factions, it seems the DC-SBM with smaller k values is a more reasonable choice.

ACKNOWLEDGMENT

The authors would like to thank Cristopher Moore and Tiago Peixoto for valuable inputs and discussions.

REFERENCES

- [1] S. Fortunato, "Community detection in graphs," *Physics Reports*, 2009.
- [2] W. W. Zachary, "An information flow model for conflict and fission in small groups," *Journal of Anthropological Research*, vol. 33, no. 4, pp. 452–473, 1977.
- [3] L. Adamic and N. Glance, "The political blogosphere and the 2004 US Election: Divided They Blog," in *Proc 3rd Intl Workshop on Link Discovery*, 2005.
- [4] S. Allesina and M. Pascual, "Food web models: a plea for groups," *Ecology letters*, vol. 12, no. 7, pp. 652–662, 2009.
- [5] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: first steps," *Social networks*, vol. 5, pp. 109–137, 1983.
- [6] S. Wasserman and C. Anderson, "Stochastic a posteriori blockmodels: Construction and assessment," *Social Networks*, vol. 9, pp. 1–36, 1987.
- [7] B. Karrer and M. E. J. Newman, *Stochastic blockmodels and community structure in networks*, Jan. 2011.
- [8] Y. Zhu, X. Yan, and C. Moore, "Oriented and degree-generated block models: generating and inferring communities with inhomogeneous degree distributions," *Journal of Complex Networks*, vol. 2, no. 1, pp. 1–18, 2014.
- [9] E. M. Airolidi, D. M. Blei, S. E. Fienberg, and E. P. Xing, "Mixed membership stochastic blockmodels," *J. Machine Learning Research*, vol. 9, pp. 1981–2014, 2008.
- [10] A. Clauset, C. Moore, and M. E. J. Newman, "Hierarchical structure and the prediction of missing links in networks," *Nature*, vol. 453, no. 7191, pp. 98–101, 2008. [Online]. Available: <http://www.nature.com/doi/10.1038/nature06830>

- [11] M. E. J. Newman and A. Clauset, "Structure and inference in annotated networks," *arXiv preprint arXiv:1507.04001*, 2015. [Online]. Available: <http://arxiv.org/abs/1507.04001>
- [12] M. Qi and G. Zhang, "An investigation of model selection criteria for neural network time series forecasting," *European Journal of Operational Research*, vol. 132, no. 3, pp. 666–680, 2001.
- [13] D. R. Anderson and K. P. Burnham, "Avoiding Pitfalls When Using Information-Theoretic Methods," *The Journal of Wildlife Management*, vol. 66, no. 3, pp. 912–918, 2002. [Online]. Available: <http://www.jstor.org/stable/3803155>
- [14] M. R. E. Symonds and A. Moussalli, "A brief guide to model selection, multimodel inference and model averaging in behavioural ecology using Akaike's information criterion," *Behavioral Ecology and Sociobiology*, vol. 65, pp. 13–21, Aug. 2010. [Online]. Available: <http://www.springerlink.com/index/10.1007/s00265-010-1037-6>
- [15] M. Handcock, A. Raftery, and J. Tantrum, "Model-based clustering for social networks," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, vol. 170, no. 2, pp. 301–354, 2007.
- [16] E. Airolidi, S. Fienberg, C. Joutard, and T. Love, "Discovery of latent patterns with hierarchical bayesian mixed-membership models and the issue of model choice," *Data mining patterns: new methods and applications*, p. 240, 2008.
- [17] P. Grünwald, *The minimum description length principle*. MIT press, 2007.
- [18] T. Peixoto, "Parsimonious module inference in large networks," *arXiv preprint arXiv:1212.4794*, 2012.
- [19] G. Claeskens and N. L. Hjort, *Model Selection and Model Averaging*. Cambridge, England: Cambridge University Press, 2008.
- [20] T. P. Peixoto, "Hierarchical Block Structures and High-Resolution Model Selection in Large Networks," *Physical Review X*, vol. 4, no. 1, p. 011047, Jan. 2014.
- [21] X. Yan, C. Shalizi, J. E. Jensen, F. Krzakala, C. Moore, L. Zdeborová, P. Zhang, and Y. Zhu, "Model selection for degree-corrected block models," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2014, no. 5, p. P05007, 2014.
- [22] K. Burnham, *Model selection and multi-model inference : a practical information-theoretic approach*, 2nd ed. New York NY: Springer, 2010.
- [23] Y. X. Wang and P. J. Bickel, "Likelihood-based model selection for stochastic block models," *arXiv preprint arXiv:1502.02069*, 2015. [Online]. Available: <http://arxiv.org/abs/1502.02069>
- [24] C. Biernacki, G. Celeux, and G. Govaert, "Assessing a Mixture Model for Clustering with the Integrated Completed Likelihood," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 7, pp. 719–725, Jul. 2000. [Online]. Available: <http://dx.doi.org/10.1109/34.865189>
- [25] R. E. Kass and A. E. Raftery, "Bayes factors," *Journal of the american statistical association*, vol. 90, no. 430, pp. 773–795, 1995.
- [26] G. Schwarz, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [27] C. Fraley and A. E. Raftery, "How many clusters? Which clustering method? Answers via model-based cluster analysis," *The computer journal*, vol. 41, no. 8, pp. 578–588, 1998.
- [28] C. Biernacki, G. Celeux, and G. Govaert, "Exact and Monte Carlo calculations of integrated likelihoods for the latent class model," *Journal of Statistical Planning and Inference*, vol. 140, no. 11, pp. 2991–3002, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378375810001631>
- [29] E. Côme and P. Latouche, "Model selection and clustering in stochastic block models with the exact integrated complete data likelihood," *ArXiv e-prints*, Mar. 2013.
- [30] M. Rosvall and C. T. Bergstrom, "An information-theoretic framework for resolving community structure in complex networks," *Proceedings of the National Academy of Sciences*, vol. 104, no. 18, p. 7327, 2007.
- [31] T. Peixoto, "Entropy of stochastic blockmodel ensembles," *PRE*, vol. 85, no. 5, p. 056122, May 2012.
- [32] T. P. Peixoto, "Model selection and hypothesis testing for large-scale network models with overlapping groups," *ArXiv e-prints*, Sep. 2014.
- [33] C. Moore, X. Yan, Y. Zhu, J. Rouquier, and T. Lane, "Active learning for node classification in assortative and disassortative networks," in *KDD '11 Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM Press, 2011, p. 841. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2020408.2020552>